Pittsburgh Series in
Philosophy and History
of Science

Series Editors:

Adolf Grünbaum
Larry Laudan
Nicholas Rescher
Wesley C. Salmon

SCIENCE AND VALUES

The Aims of Science and Their Role in Scientific Debate

Larry Laudan

University of California Press
Berkeley Los Angeles London

CONTENTS

Acknowledgments ix
Preface xi

One Two Puzzles about Science: Reflections on Some Crises in Philosophy and Sociology of Science 1

The Consensual View and the Puzzle of Agreement 3
The "New Wave" Preoccupation with Dissensus 13

Two The Hierarchical Structure of Scientific Debates 23

Factual Consensus Formation 26
Methodological Consensus Formation 33

Three Closing the Evaluative Circle: Resolving Disagreements about Cognitive Values 42

The Covariance Fallacy 43
The Reticulated Model and the Mechanics of Goal Evaluation 50
The Reticulated Model of Scientific Rationality 62

Four Dissecting the Holist Picture of Scientific Change 67

Kuhn on the Units of Scientific Change 68
Kuhn's Critique of Methodology 87

Five A Reticulational Critique of Realist Axiology and Methodology 103

Epilogue 138
References 141
Index 145

every bit as daunting as those that proved to be the undoing of empiricist methodology and Mertonian sociology. More specifically, many recent theorists, while labeling classical philosophy and sociology as impoverished, have ignored the central issues with which their predecessors were grappling. We can scarcely claim to have moved significantly beyond the work of the 1940s and 1950s unless we can make some sense of the striking facts that scholars of that generation rightly regarded as basic features of science. We either have to deny with Feyerabend that rational scientists could ever exhibit widespread agreement (and that seems to run counter to the record), or else we have to find some account of dissensus which is not so robust that it precludes the very possibility of frequent and widespread agreement. Until we manage to account for a Janus-faced science, we cannot seriously claim to have understood what we are about.

This book is an attempt to move us some steps forward in that direction. In succeeding chapters I focus chiefly on describing the various levels at which scientific disagreement can occur. In each instance we will be exploring how far one can expect disagreements to be amenable to rational analysis and rational closure. As we shall come to see, the full-blown Leibnizian ideal cannot be plausibly resurrected, for there remain many scientific controversies that cannot be rationally terminated, even with the best will in the world. On the other hand, we will discover a very large range of cases where there is appropriate analytic machinery for understanding how many scientific controversies can be brought to a reasonably definitive resolution.

Chapter Two

THE HIERARCHICAL STRUCTURE OF SCIENTIFIC DEBATES

In any community as diverse as the scientific one, and especially in one with such a deeply entrenched tradition of challenges to authority, where successful breaks with tradition are handsomely rewarded, consensus is not born but made. Because agreement typically emerges out of prior disagreement, it is useful to cast the puzzle of consensus formation in this form: How is it that a very high proportion of scientists, who previously had different (and often mutually incompatible) views about a particular subject, can eventually come to hold substantially identical views about that subject? Put this way, the problem of consensus formation is a problem about the dynamics of convergent belief change.

The best-known contemporary solution to the problem of consensus formation in science involves postulating what I call the hierarchical model of justification, although it is perhaps more commonly known as the theory of instrumental rationality. Proponents of this model[1] typically envisage three interrelated levels at which, and by means of which, scientific consensus is forged. At the lowest level of this hierarchy are disputes about matters of fact. By the phrase "matters of fact" I mean to refer not only to assertions about directly observable events but to all manner of claims about what there is in the world, including claims about theoretical or unobservable entities. For obvi-

1. Among the influential philosophical advocates of the hierarchical model are Karl Popper, C. G. Hempel, and Hans Reichenbach.

ous reasons, I call debates of this sort "factual disagreements," and agreement at this level, "factual consensus." According to the standard account, scientists resolve factual disagreements and thus forge factual consensus by moving one step up the hierarchy to the level of shared methodological rules. The rules may be mechanical algorithms for generating factual statements. But, much more typically, the rules will simply be constraints or injunctions concerning the attributes we should seek (e.g., independent testability) or avoid (e.g., ad hocness) in our theories. As normally understood, such rules, which are basically principles of empirical support and of comparative theory assessment, provide directives for ascertaining, at least in a qualitative sense, how much support (i.e., confirmation or disconfirmation) the available evidence provides for the theories under evaluation. If two scientists disagree about whether one rival factual claim or another is more worthy of belief, they have (in this view) but to compare the weights of support enjoyed by their respective claims in order to terminate their disagreement.

According to this model, decisions between competing theories may be likened to the way in which our courts settle, or are supposed to settle, civil contests: relevant evidence is presented; the court agrees to shape its verdict according to well-established jurisprudential rules of evidence; an impartial verdict is "guaranteed" because the issue is settled in light of the rules rather than of personalities; and, finally, all the parties to the case agree to abide by the verdict. In the same way, the hierarchical model requires that scientists submit their factual disputes to a kind of invisible "science court" (in this case, the practitioners of one or another scientific specialty). The "scientific jury" is expected to make its choice according to rules of evidential support agreed to by all scientists working in that specialty. Such prior agreement is thought to guarantee that the "verdict" will be both impartial and acceptable to all parties. At first glance, this approach has much going for it. It can explain not only why scientific disagreements often issue in consensus, but also the rapidity with which they do so. (And as pointed out in chapter 1, the really remarkable thing about many scientific controversies is how quickly they are brought to a definitive resolution.)

The thesis that factual disagreements can be resolved by invoking the relevant rules of evidence is, of course, just a modern exemplification of what I have earlier called the "Leibnizian ideal." But whereas

that ideal in its original form imagined that all factual disagreements could be terminated by invoking the relevant rules, latter-day proponents of methodological rules tend to be more modest. They continue to believe that some disagreements can be immediately resolved by utilizing the available evidence (and the shared rules). Failing that, however, they go on to say that the rules are often sufficiently specific to indicate procedures for the collection of such additional evidence as will bring the issue to a definitive resolution. The rules themselves vary from the highly general ("formulate testable and simple hypotheses") to those of intermediate generality ("prefer the results of double-blind to single-blind experiments"), to those specific to a particular discipline or even subdiscipline ("make sure to calibrate instrument x against standard y"). To the extent that these procedural or methodological rules are accepted by all parties to the dispute, and insofar as they are sufficiently specific to determine a choice between the available rivals, they should indeed suffice for the mediation of factual controversies. And a staggeringly large proportion of factual disputes have evidently been ended simply by observing the relevant methodological procedures.

Sometimes, however, scientists disagree about the appropriate rules of evidence or procedure, or about how those rules are to be applied to the case at hand.[2] In such circumstances, the rules can no longer be treated as an unproblematic instrument for resolving factual disagreement. When this happens, it becomes clear that a particular factual disagreement betokens a deeper methodological disagreement. In the standard hierarchical view, such methodological controversies are to be resolved by moving one step up the hierarchy, that is, by reference to the shared aims or goals of science. This suggestion is a natural one

2. The fact that scientists sometimes subscribe to different methodological principles stands as an important anomaly to Andrew Lugg's otherwise very stimulating account (1978) of the causes of scientific disagreement. Lugg's general thesis is that disagreements about substantive scientific matters may arise because scientists have, in effect, differential degrees of access to the data base potentially relevant to the assessment of any theory. Thus, practitioners working in one subspecialty will usually have a different perspective on the degree of support enjoyed by competing theories from that employed by the corresponding specialists in another branch of the science. Lugg is surely right in deeming difference as one source of disagreement. But he is wrong in suggesting that a community of "ideal" observers, each of whom had access to all the relevant available evidence, would be bound to produce a consensual science. For to say as much is to ignore the real differences of methodological orientation which scientists, even when confronted by the same evidence, often exhibit.

to make, for a little reflection makes clear that methodological rules possess what force they have because they are believed to be instruments or means for achieving the aims of science. More generally, both in science and elsewhere, we adopt the procedural and evaluative rules we do because we hold them to be optimal techniques for realizing our cognitive goals or utilities. Hence, when two scientists find themselves espousing different and conflicting methodological rules (and assuming, as the standard account does, that they have the same basic aims), they can in principle terminate their disagreement at the methodological level by determining which of the rival rules conduce(s) most effectively to achieving the collective goals of science. I call this third stage, where basic cognitive aims are involved, the axiological level.

We can therefore sum up the prevailing philosophical view on scientific disagreements succinctly: disagreements about factual matters are to be resolved at the methodological level; methodological differences are to be ironed out at the axiological level.[3] Axiological differences are thought to be either nonexistent (on the grounds that scientists are presumed to share the same goals) or else, should they exist, irresolvable (see fig. 1). In the rest of this chapter I explore some of the strengths of, and objections to, this pervasive model of consensus formation.

FACTUAL CONSENSUS FORMATION

One apparent weakness of the hierarchical view lies in its central assumption that methodological rules will, at least in principle, always pick out one factual claim, to the exclusion of all its possible rivals, as uniquely supported by those rules. Yet it is notorious that methodological rules usually underdetermine a choice among factual claims in the sense that, although the rules plus the available evidence will exclude many factual claims or hypotheses, a plethora of possible hypotheses

3. There is an excellent formulation of the hierarchical view of the relation between methods and aims in Gutting, 1973. Squarely in the hierarchical tradition, Gutting stresses that methodological disputes can often be resolved by looking to the "founding intentions" or the aims of the members of a scientific community. This is to ignore (*a*) that those intentions themselves do not remain constant through time and, more important, (*b*) that scientists are often unable to agree on what the founding intentions of a discipline are. In such circumstances, Gutting, like every other advocate of the hierarchical view, proposes no remedies for mediating disagreement (see esp. Gutting, 1973, pp. 277 ff.).

Level of Disagreement	*Level of Resolution*
Factual	Methodological
Methodological	Axiological
Axiological	None

Fig. 1. The Simple Hierarchical Model of Rational Consensus Formation

often remains methodologically admissible. Among the admissible hypotheses may be some that are evidentially equivalent, in that no conceivable evidence could ever discriminate among them. It is often said, for instance, that certain versions of wave and matrix mechanics are observationally equivalent. In such circumstance, it is clear that no observations could ever decisively choose between them. In other cases the admissible hypotheses will be evidentially distinct (in the sense that there is conceivable evidence that would differentially support them) but such that the existing evidence and the prevailing rules do not provide grounds for a preference. Since the set of factual claims supported by certain methodological rules can be shown always to include several contraries of one another,[4] it is often charged that the underdetermination of theory by the relevant rules and evidence makes a mockery of the hierarchical model of consensus formation (at least when that model is applied to the adjudication of disagreements about theories). But this commonplace criticism generally misses the point. What we usually want to explain in any particular instance of consensus formation is not how scientists were able to agree to accept a certain hypothesis rather than every other possible hypothesis. If that were our explanatory puzzle, then methodological underdetermination would indeed preclude an answer. Equally, if the partisans of rival scientific theories were looking to find out if one of their theories was better supported than all possible rivals, then this sort of underdetermination would be debilitating. But this is to misstate the problem in both cases. What scientists are trying to decide (and here our earlier judicial analogy is especially apt) is not whether their theories will last for all time, or whether they will always stack up favorably against all conceivable rivals; rather they are trying to decide which of the theories presently contending in the scientific marketplace is best supported by the evidence. Scientists, in the view I am advocating, should be seen, not as looking for simply the best theory, but rather for the best theory they

4. For a more detailed treatment of underdetermination, see chap. 4 below.

can find. I am suggesting that a more constructive, and a more realistic, way of formulating the problem of consensus formation is this: Given that some scientists once believed one theory and that others once accepted a rival to it, why do they all now accept the latter?[5] In other words, we are usually seeking to understand why a particular preference was made consensually to endorse one from among a finite (and usually quite narrow) range of articulated rival theories. To make it more concrete, we want to find answers to such questions as these: Why did geologists, who once heaped such scorn on the theory of continental drift, eventually come to accept plate tectonics? Why did physicists, who once were advocates of the corpuscular theory of light, eventually come to believe that the undular theory was a better bet?

Once we put the problem of consensus formation in this comparative mode, we take much of the sting out of the argument from underdetermination. For in many instances the shared rules and the existing evidence neatly partition the extant theories under active consideration into two sets, those that, according to the rules, are supported by the evidence and those that are not. In the event that both parties to a controversy draw their theories from among the elements of the former set, then the methodological rules will clearly be insufficient to mandate a preference and the participants will simply have to agree to disagree—pending the accumulation of further evidence. Instances of such (temporary) underdetermination arise fairly often, and they are, of course, the most interesting ones for historians and sociologists of science to examine. But the appeal of such long-term standoffs ought not blind us to the fact that they are arguably a tiny minority of cases. Most of the time all the parties to the dispute will agree that there comes a point where the rules unambiguously warrant a preference among the extant competing theories.

One might wonder how I can say that rules can simultaneously underdetermine belief and determine a preference. This distinction is elementary, but it needs to be attended to carefully, for numerous authors have been confused by it. Essentially, a rule (or a set of rules) partitions prospective beliefs into two classes, the permissible and the impermissible. A belief is permissible precisely when, among the various alternatives to it under active consideration, none has a higher

5. I have formulated the problem for a choice between two rival theories; it could readily be generalized to apply to a larger number.

degree of empirical support than it does.[6] If we define permissible beliefs in this way, we see that several rival claims about a certain matter of fact may all be permissible in the face of a certain body of evidence. That is, they may all be equally well supported. However, and here is the contrast with strict underdetermination, there can be circumstances in which one rival alone among those under active consideration is significantly better supported by the evidence than the others. In such cases that rival alone is permissible; its acceptance, rather than the acceptance of any of its known rivals, is unambiguously dictated by the germane rules and evidence.

Of course, the preferred rival may still be underdetermined in the strict sense of the term, since there may be conceivable, but thus far unconceived, rival theories that would be as well supported as our permissible choice. The crucial point here is that even when a rule underdetermines choice in the abstract, that same rule may still unambiguously dictate a comparative preference among extant alternatives. It will do so specifically when we are confronted with a choice between (in the simplest case) two candidate theories, one of which is (methodologically) permissible and the other not.

For instance, the rules and evidence of biology, although they do not establish the unique correctness of evolutionary theory, do exclude numerous creationist hypotheses—for example, the claim that the earth is between 10,000 and 20,000 years old—from the permissible realm and thus provide a warrant for a rational preference for evolutionary over creationist biology. If we once grant that theory appraisal is a comparative matter, that the scientist is generally making comparative judgments of adequacy among available rivals rather than absolute judgments about the best possible theory, then it becomes clear that comparative preferences may be warranted even when the selection of the best possible theory is beyond our justificatory resources.

I review this familiar terrain because several recent writers[7] have been apt to ignore the difference between (what I am here calling)

6. Tough-minded readers may want to work with a more demanding sense of permissible belief, according to which a belief is regarded as warranted only if it enjoys stronger support from the rules than all its extant rivals. (Otherwise, we might find ourselves in a situation where contrary beliefs were all permissible.) I would myself resist the stronger characterization of "permissible," but the argument I develop here in the text will work, whichever formulation of permissible one inclines to.

7. Including Quine, Hesse, and Bloor.

choices and preferences. Specifically, they appear to have argued that, since theory choice is underdetermined by methodological rules, it follows that no rational preference is possible among rival theories, which entails, in turn, that every theory is as well supported as any other, and that every party to a scientific debate is thus as rational as every other. Our discussion should have made clear that the thesis of underdetermination, in its viable forms, provides no justification for the view that all theories are equally well confirmed or equally rational in light of a given body of empirical information. The implications of this point might be discussed at greater length if our concern were chiefly with the logic and methodology of theory evaluation. But it is sufficient for our purposes here to note that the slide from underdetermination to what we might call cognitive egalitarianism (i.e., the thesis that all beliefs are epistemically or evidentially equal in terms of their support) must be resisted, for it confusedly takes the fact that our rules fail to pick out a unique theory that satisfies them as a warrant for asserting the inability of the rules to make any discriminations whatever with respect to which theories are or are not in compliance with them.

In its extreme form, such egalitarianism amounts to a radical version of epistemic relativism, one formulation of which can be found in the thesis of the sociologist Harry Collins that, because our beliefs are not uniquely determined by the evidence, our beliefs can reasonably be presumed to be independent of the evidence for them! In a less extreme and more carefully articulated form, one sees the same sort of argument in Thomas Kuhn's recent work (especially in *The Essential Tension*) on the role of rules and values in science. Like Wittgenstein before him, Kuhn makes much of the fact that there are no algorithms for scientific choice; that is, we have no agreed-upon rules that would uniquely pick out the best possible theory in light of a given body of evidence and background assumptions. What rules we do have, Kuhn seems to think, are invariably so ambiguous that they are of little practical use in making specific theory choices. Kuhn is not dismissive of rules; nor, like certain skeptical sociologists, does he regard them as so many post hoc rationalizations. To the contrary, Kuhn insists that methodological rules play both a causal and a justificatory role in scientific belief. But because Kuhn also believes that the rules of appraisal are highly ambiguous and fluid (he likens them to such homespun and largely uninformative maxims as "look before you leap"), he inclines to assign them a minimal role. He says, for instance, that because we

have no rules that would unambiguously pick out a single theory to the exclusion of all other possible theories about the relevant domain, it is inevitable that the choice between any two theories could always go one way or the other, given any set of values or norms about what we expect our theories to achieve. Kuhn is not denying that rules play a role in the choice of scientific theories, but he is insisting that their intrinsic ambiguity precludes the possibility of decisive preferences ever being justified on the basis of shared methodological rules. Kuhn grants that "such canons [his examples are accuracy, simplicity, generality, etc.] do exist and should be discoverable," but he goes on to insist that "they are not by themselves sufficient to determine the decisions of individual scientists."[8] If all that Kuhn is saying here is that general rules and values underdetermine choice (in the strict sense defined above), one might not quarrel with his claim; but, as his subsequent discussion makes clear, he is asserting—to use my earlier language—that preferences are underdetermined as well. He puts it this way: "every individual choice between competing theories depends on a mixture of objective and subjective factors, of shared and individual criteria."[9] Kuhn thinks that this state of affairs must be so because, according to his analysis, the "objective" criteria shared by scientists cannot justify one preference to the exclusion of another. Since, as Kuhn knows perfectly well, scientists do in fact voice theory preferences, he takes this as evidence that they must be working with various individual and idiosyncratic criteria that go well beyond the shared ones. Without the latter, he seems to say, how could scientists ever have preferences? But neither Kuhn nor anyone else has shown, either in fact or in principle, that such rules and evaluative criteria as are shared among scientists are generally or invariably insufficient to indicate unambiguous grounds for preference of certain theories over others.

Indeed, using some of Kuhn's own criteria, we can show precisely how a rationally based preference can arise. Suppose that a scientist is confronted with a choice between specific versions of Aristotle's physics and Newton's physics. Suppose moreover that the scientist is committed to observational accuracy as a primary value. Even granting with Kuhn that "accuracy" is usually not precisely defined, and even

8. Kuhn, 1977, pp. 324-325.
9. Ibid., p. 325.

though different scientists may interpret accuracy in subtly different ways, I submit that it was incontestable by the late seventeenth century that Newton's theory was empirically more accurate than Aristotle's. Indeed, even Newton's most outspoken critics conceded that his theory was empirically more accurate than all its ancient predecessors. Similarly, if it comes to a choice between Kepler's laws and Newton's planetary astronomy—and it does come to such a choice since the two are formally incompatible—and if our primary standard is, say, scope or generality of application (another of the cognitive values cited by Kuhn), then our preference is once again dictated by our values. At best, Kepler's laws apply only to large planetary masses; Newton's theory applies to all masses whatsoever. Under such circumstances the rule, "prefer theories of greater generality," gives unequivocal advice. There will, of course, be times when the shared criteria are too ambiguous to yield a definitive preference; such cases have tended to preoccupy Kuhn in his analysis of scientific change. But when he generalizes from such arguably idiosyncratic cases to make claims about "every individual choice between competing theories" being immune to resolution by the shared rules, and when he goes on to suggest that no one has ever denied that his account is "simply descriptive of what goes on in the sciences at times of theory choice,"[10] it is necessary to set the record straight by pointing out that neither Kuhn nor anyone else has shown that most (let alone all) theory choice situations exhibit the impotence of "shared criteria" to determine a preference.

In other words, although Kuhn is surely right to stress that we have no perfectly general logic of confirmation or comprehensive theory of evidence, just as he is right to say that many rules are ambiguous, he is wrong to see these facts as sustaining the claim that the application of shared scientific rules or values to a specific choice situation will always be (or always has been) ambiguous or unavailing. There are some situations where it is unavailing; we have already discussed some of these under the general head of "underdetermination." But Kuhn has neither shown, nor even made plausible, the claim that the rules, evaluative criteria, and values to which scientists subscribe are so ambiguous in application that virtually any theory can be shown to satisfy them. And we must constantly bear in mind the point, already made, that even when theories are underdetermined by a set of rules, there will

10. Ibid., p. 325.

typically be many theories ruled out by the operant rules; and if one party to a scientific debate happens to be pushing for a theory that can be shown to violate those rules, then the rules will eliminate that theory from contention.

It follows from this analysis that the hierarchical model is not (as many of its critics claim) rendered toothless by the argument from underdetermination. To the contrary, it seems entirely reasonable to say that many disputes about matters of fact have been terminated by invoking shared procedural rules. But it is important to realize what "hedges" are built into this reformulation of the hierarchical view. Not all factual disagreements can be resolved in this Leibnizian fashion because two or more extant rivals may be equally well supported by the existing rules and evidence. Similarly, it may happen (a case that I will explore shortly) that scientists differ about which evidential rules should be applied to the case at hand. The fact that this version of the hierarchical model leads us to expect consensus only in some instances will satisfy neither the archrationalist advocates of the Leibnizian ideal (who want an in-principle, instant termination for every factual disagreement) nor the proponents of radical underdetermination (who believe in the potential prolongation of every disagreement indefinitely). But the strength of this version of the hierarchical model is that it can (unlike the underdeterminationists) specify circumstances in which we would expect a factual disagreement to dissolve into consensus and it can (unlike the Leibnizian idealists) also specify a broad range of circumstances in which we would expect factual dissensus to endure.

METHODOLOGICAL CONSENSUS FORMATION

Thus far we have been uncritically accepting one central feature of the Leibnizian picture, even while rejecting others, for we have been discussing situations in which a group of scientists shares a set of cognitive aims and methodological rules, but disagree about certain factual matters. It sometimes happens that disagreements go deeper than this. We occasionally see scientists unable to agree even about the appropriate methodological and procedural rules to bring to bear on the choice of hypotheses or theories. One scientist, for instance, may believe (with Popper) that a theory must make surprising, even startling, predictions, which turn out to be correct, before it is reasonable to accept it.

Another may be willing to accept a hypothesis so long as it explains a broad range of phenomena, even if it has not made startling predictions. A third may say (with Nagel) that no theory is worth its salt until it has been tested against a wide variety of different kinds of supporting instances. A fourth may believe that a very large number of confirmations is probatively significant, regardless of the variety they exhibit. A fifth may demand that there be some direct and independent evidence for the existence of the entities postulated by a hypothesis before it can reasonably be accepted. All these familiar methodological principles of theory acceptance are at odds with one another, and each has found prominent advocates in recent science and philosophy.

What are we to say when scientists disagree about the rules of the scientific game? Is there any hope, apart from artificially imposing closure from outside the system, that scientists can reasonably resolve their methodological differences? Or are they caught up in a nasty normative incommensurability which defies rational solution? The answer to that question depends, according to the hierarchical model, on our ability to resolve such methodological controversies by moving one step up the cognitive ladder of justification.

To see how the hierarchy comes into it, we have to remind ourselves of the function of, and thus of the rationale for, rules in general. We generally agree to govern an activity (whether it is science, chess, or parliamentary debate) by a particular set of rules when we think that those rules will enable us to achieve the ends or goals that define the telos of the activity. Scientists presumably have the methodological rules they do because they suppose that following the rules in question will bring about, or at least bring them closer to, the realization of their cognitive or doxastic aims. So conceived, methodological rules are nothing but putative instrumentalities for the realization of one's cognitive ends; in a word, the rules of science are designed simply as means to cognitive ends or tools for performing a task. To think about rules in this way immediately suggests an answer to our question about how to resolve disputes about rules and, in the process, offers a way out of the incommensurability that apparently ensues when people cannot agree to play the game by the same rules. It should now be clear that if two scientists disagree about the appropriate rules but agree about some "higher" cognitive values or ends, then we can possibly resolve the disagreement about rules by ascertaining which set of rules is most likely to realize the common cognitive aims. Once we know the answer

to that question, we will know what the appropriate methodological rule(s) should be, and we will have a resolution of the methodological disagreement, at least insofar as it was rationally founded in the shared axiology.

But as plausible as this answer is in general terms (and I believe that it is on the right track), it makes things seem a good deal tidier than they really are and it leaves several important issues unaddressed. For instance, it presupposes that a given set of cognitive aims or goals is uniquely associated with a certain set of methodological rules, that is, that a particular set of rules can be shown to be the only optimal means for achieving the values in question. It is generally difficult, and in certain cases patently impossible, to exhibit that a particular set of rules is the best possible way for realizing a certain set of values. More modestly, but just as usefully, we may be able to show that following a particular set of rules will indeed realize a certain epistemic value, although it may be quite another matter to show that these rules are the best or the only means for achieving the values we desire. And if we cannot demonstrate the latter, then we cannot argue for the blanket superiority of those rules over all their conceivable rivals for realizing the values in question. In short—and this is crucial—cognitive aims typically underdetermine methodological rules in precisely the same way that methodological rules characteristically underdetermine factual choices. The classic philosopher's quest to show that certain procedures of inquiry are the only route to a desired epistemic end is a largely misguided one, because we cannot usually enumerate, let alone examine, all the possible ways of achieving a certain goal. And without such an examination we usually have no warrant for asserting that one set of methods is superior to all the others. To make the point less abstractly, one can observe that with respect to such familiar cognitive goals as truth, coherence, simplicity, and predictive fertility, scholars have not managed to show that there is any set of rules of empirical investigation which uniquely conduces to their realization. In sum, rules, like theories, are underdetermined by the relevant constraints.

And a good thing they are too. For if there was only one set of rules for realizing any specific set of cognitive aims, we should have to conclude that it was irrational for scientists sharing the same ends or values ever to disagree about the appropriate rules for implementing those values. Yet such disputes are chronic in the history of science and philosophy. Consider, for instance, the 150-year-long (and still on-going) controversy about the so-called rule of predesignation.[11] This rule specifies that a hypothesis is tested only by the new predictions drawn from it, not by its ability post hoc to explain what was already known. A host of prominent thinkers have been arrayed on each side of this issue (Whewell, Peirce, and Popper for predesignation; Mill and Keynes, among others, against it). All parties to the controversy would, I believe, subscribe to substantially the same cognitive aims. They seek theories that are true, general, simple, and explanatory. Yet no one has been able to show whether the rule of predesignation is the best, or even an appropriate, means for reaching those ends. That failure is entirely typical. There is almost no cognitive value and associated methodological rule which have been shown to stand in this one-to-one relation to each other. So far as we know, there may be equally viable methods for achieving all the cognitive goals usually associated with science. It is for just this reason, incidentally, that the much sought-after "scientific method" may be a will-o'-the-wisp. To seek the rules of *the* scientific method is to presuppose that there is only one legitimate means to the achievement of the shared cognitive aims of science. Because there may well be a variety of methodological rules which conduce equally well to achieving our cognitive values, it follows that the coexistence of nonidentical methods of inquiry may well be a permanent feature of scientific life.[12]

But, the skeptic may ask, if we can never justify a methodological rule by showing it to be the only or the optimal means for achieving a certain cognitive value, how can we—as I have claimed—use the aims and values of science as a tool for resolving disputes about methodological rules? Shared goals can often mediate controversies about rules precisely because they impose constraints on the class of permissible rules. We can sometimes show, for instance, that certain rules are nonconducive, even counterconducive, to achieving certain values. If, for example, one of our cognitive values is generality and breadth of scope in our theories, then it is quite clear that any appraisal rule that

11. I discuss the early history of the rule of predesignation in Laudan, 1981.

12. A second, probably insuperable, obstacle to the quest for *the* scientific method is the absence of a full consensus among scientists about what the cognitive goals of science should be. Without that agreement we should scarcely be surprised if some dissensus about methods persists. (But it should be stressed, as noted below, pp. 43 ff., that a lack of consensus about cognitive aims does not necessarily entail the existence of any dissensus about the appropriate methods.)

insists that we should accept only highly probably theories is unsatisfactory. Precisely because one can (following Popper) demonstrate an incompatibility between high probability and broad explanatory scope, we can quickly exclude a rule urging maximally high probabilities from the methodological repertoire of anyone whose primary cognitive aims include generality. Hence, one role of aims in resolving methodological disagreements is to eliminate certain methodological rules, by virtue of their irreconcilability with those aims.

Besides this exclusionary function, aims can sometimes play a positive role, as, for example, when we can show that certain methodological rules promote certain cognitive values. For, although we generally cannot show that a particular rule is the best of all possible rules for realizing a certain end, we often can show both (1) that it is one way of realizing a particular end, and (2) that it is better than all its rivals that are under active consideration. And this result will often be sufficient to resolve a particular disagreement about methodological rules. If, for instance, we can show that one particular rule is a better way to achieve given cognitive ends than a rival rule is, and if the controversy is specifically between advocates of those two rules (who, moreover, have a shared set of cognitive aims), then we are in a clear position to move the controversy rationally to closure.

But although the invocation of shared goals may sometimes make methodological consensus possible, it is crucial to stress that this is not a cure-all for all manner of methodological disagreements. It may happen, for instance, that both parties to the controversy are advocating methodological rules that are, so far as we can see, equally effective means to achieving the cognitive values in question. Even more seriously (and more commonly), we may find ourselves in a situation where we endorse a broad range of cognitive aims or values simultaneously (say, simplicity, coherence, and empirical accuracy). A methodological rule may tend to promote the realization of one of these values while it thwarts the realization of another. Is the rule, under such circumstances, a good one or a bad one? The answer depends, in part, on how we respectively weight the goals we endorse. If the rule promotes those goals that are the most important to us, we may find the rule acceptable. Yet someone who subscribes to the same cognitive goals, but weights them differently, may well find the rule unacceptable. Methodological disagreements such as these show little promise of adjudication by bringing aims explicitly into play. Equally problematic is the situation where the disagreement about rules derives from an even deeper disagreement about cognitive aims (discussed in detail in chap. 3).

But before we move on to deal with the very tricky case of disagreements about basic goals, we should pause a moment to reflect on some of the implications of our account of the resolution of methodological disagreements. If the analysis sketched here is correct, we have to realize that the vertical hierarchy of facts-rules-aims has a rather more complex structure than figure 1 and our earlier discussion might suggest. Although we appraise methodological rules by asking whether they conduce to cognitive ends (suggesting movement up the justificatory hierarchy), the factors that settle the question are often drawn from a lower level in the hierarchy, specifically from the level of factual inquiry. Factual information comes to play a role in the assessment of methodological claims precisely because we are continuously learning new things about the world and about ourselves as observers of that world. Such knowledge comes to be formulated in theories concerning, among other things, the various complexities of the process of collecting evidence. To take an elementary example, we have learned that nature does not offer information to us in a random or statistically representative way. As medium-sized objects in a world replete with the very small and the very large, the entities and processes we are most likely to encounter in our everyday scrutiny of the world are highly unrepresentative of that world in many crucial respects. Once we learned that fact, it became necessary to develop elaborate sampling techniques in order to make our evidence more representative than it would have been had we simply collected whatever information casually came our way. Contrary to popular conception, the superiority of randomly collected data over nonrandomly collected data is not a discovery made by mathematics or formal logic as a feature of inquiry in all possible worlds. It is because the particular world we inhabit turns out to be so uncooperative and, as we have learned that only through hard-won experience, we find it appropriate to insist on a variety of stringent rules about sampling in order to achieve our cognitive ends.

Consider a different example. Within the past fifty years we have learned a great deal about the so-called placebo effect. In brief, many patients report an improvement upon being given any apparent medication, even if (unbeknownst to the patient) it is pharmacologically inert. Until we learned this, simple controlled experiments were regarded as sufficient tests of therapeutic efficacy. To find out if a drug worked, scientists formerly would give it to one group of patients and nothing whatever would be administered to a second (control) group with comparable symptoms. If the former reported higher rates of improvement than the latter, the test was taken as good evidence that the drug was effective. But once we learned about the placebo effect, it became clear that the testing of various new therapies (including virtually all drugs) had to be rather more complicated than we had once imagined. Specifically, scientists realized that they needed to resort to so-called single-blind tests, which make appropriate allowances for spurious reports of efficacy (i.e., reports of efficacy based largely on the patient's expectations of betterment). To make matters still more complex, once we learned further that those administering drug tests could often, if unconsciously, transmit their own therapeutic expectations to the patients they were examining, it became necessary to formulate techniques for double-blind tests of medications and other therapies.

In all these instances our views about the proper procedures for investigating the world have been significantly affected by our shifting beliefs about how the world works.[13] Factual beliefs thus shape methodological attitudes, every bit as much as our goals do. A priori reflection on our cognitive aims would not have necessitated the development of randomizing techniques or of single- and double-blind experiments. It was, rather, the realization, given what we had learned of the world, that our cognitive aims could be achieved more effectively by resorting to some such techniques which mandated their introduction and provided their justification.

What the foregoing suggests is that the methodology and epistemology of science, whose central concern is with the assessment of various rules of inquiry and validation, should be conceived, far more than they normally are, as empirical disciplines.[14] Crediting or discrediting

13. There is thus a central, but nonvicious, circularity about our evaluative procedures: we use certain methods for studying the world and those very methods may serve initially to authenticate discoveries that expose the weaknesses of those selfsame methods.

14. This suggestion should not be confused with Quine's familiar argument about "naturalized epistemology." Whereas Quine sees close affinities between psychology and epistemology, I am not committed to that particular program of reduction. My insistence, rather, is that the appraisal of proposed cognitive methods and aims requires extensive empirical research. That research will often have nothing specifically

a methodological rule requires us to ask ourselves whether the universe we inhabit is one in which our cognitive ends can in fact be furthered by following this rule rather than that. Such questions cannot be answered a priori; they are empirical matters. It follows that scientific methodology is itself an empirical discipline which cannot dispense with the very methods of inquiry whose validity it investigates. Armchair methodology is as ill-conceived as armchair chemistry or physics. There is nothing particularly new in this conception of the character of scientific methodology. More than half a century ago, John Dewey—ever eager to naturalize the a priori—repudiated that conception of methodology which saw it as "an affair . . . of fixed first principles . . . of what the Neo-scholastics called criteriology."[15] What is remarkable is that, despite massive evidence to the contrary, some philosophers of science still tend to imagine that methodology is protoscientific, that is, prior to and independent of the kind of empirical inquiry which science itself represents.

Unfortunately, many of those authors who have recently argued that epistemology and methodology should be rendered more empirical or more "naturalistic" seem to have assumed that a naturalizing of epistemology would necessarily empty it of normative force, presumably on the grounds that a truly empirical epistemology would be exclusively descriptive. Granting for a moment that there is no hard-and-fast line between normative and descriptive activities, the presumption that an empirical theory of knowledge would be void of normative claims is nonsense. Once we realize (as this chapter should make clear) that methodological norms and rules assert empirically testable relations between ends and means, it should become clear that epistemic norms, construed of course as conditional imperatives (conditional relative to a given set of aims), should form the core of a naturalistic theory of scientific knowledge. When authors like Barry Barnes suggest that "the normative concerns common among epistemologists are difficult to reconcile with an empirical orientation to science,"[16] they reveal how deeply they have misunderstood the fact that an empirical approach to epistemology requires attention to precisely those normative linkages between cognitive ends and means which constitute scientific rationality.

To conclude, we find ourselves with mixed results. There do seem to be circumstances under which both factual and methodological disagreements can be brought to a rational resolution by seeking shared assumptions at a higher level. The familiar idea that agreement about factual matters is possible only among those who already accept the same methodological rules is clearly seen to be too restrictive, for it fails to reckon with the fact that, by moving a step up the hierarchy to aims and goals, dissensus about rules may dissolve into consensus. And that consensus about rules, once in place, may be sufficient to resolve the disagreement about the contested factual matters. But we have also discovered a number of problem cases. When the shared rules fail to dictate a factual preference, when the shared goals fail to specify a methodological preference, when values are shared but not weighted equally, and when values are not fully shared, we seem to be confronted by an irresolvable disagreement—irresolvable, that is, if we stick to the limited resources of the classical hierarchical model. Most crucial of all is the situation where scientists subscribe to different goals. Such circumstances occur sufficiently often to pose a fundamental challenge to the very idea that science is a rational and progressive undertaking. But, as we shall see in chapter 3, that challenge is not always so stark as it may seem.

---

to do with psychological phenomena. Whereas Quine evidently sees naturalized epistemology as constituting a subfield within psychology, I argue that an empirical epistemology or methodology is neither a part of, nor subordinate to, psychology. And typically, it would draw more heavily on physics and biology than on psychology.

15. Dewey, 1938.

16. Barnes, 1982, p. 63.